2007 IEEE International Conference on
Robotics and Automation
Roma, Italy, 10-14 April 2007



# Homography Based Kalman Filter for Mosaic Building. Applications to UAV position estimation

Fernando Caballero*, Luis Merino†, Joaquín Ferruz* and Aníbal Ollero*
{caba,merino,ferruz,aollero}@cartuja.us.es
Robotics, Vision and Control Group
*University of Seville, Camino de los Descubrimientos s/n, 41092, Sevilla, Spain
†Pablo de Olavide University, Crta. Utrera km1, 41013, Sevilla, Spain

*Abstract*— **This paper presents a probabilistic framework where uncertainties can be considered in the mosaic building process. It is shown how can be used as an environment representation for an aerial robot. The inter-image relations are modeled by homographies. The paper shows a robust method to compute them in the case of quasi-planar scenes, and also how to estimate the uncertainties in these local image relations. Moreover, the paper describes how, when a loop is present in the sequence of images, the accumulated drift can be compensated and propagated to the rest of the mosaic. In addition, the relations among images in the mosaic can be used, under certain assumptions, to localize the robot.**

## I. INTRODUCTION

Knowledge about the environment is a critical issue for autonomous operations. In general, the environment representation depends on the kind of sensors used to estimate it and the tasks and circumstances in which the vehicle will be involved. This paper addresses the environment representation problem for aerial vehicles and its estimation by means of monocular imagery.

In the case of an aerial robot that is not affected by obstacles at the flight altitude, geo-referenced mosaics can be sufficient as environment model for certain tasks. A mosaic is built by aligning to a common frame a set of images gathered while the aerial vehicles is moving. Image alignment is based on feature matching techniques and projective geometry tools.

If the scene is planar (which is, in general, a valid approximation when the ratio between the distance to the scene and the ground elevation is high [1], [2]), a set of matches between two views can be used to estimate a homographic model for the apparent image motion. This model can be composed with the estimated motion of the rest of images in order to build the mosaic. Moreover, for a planar scene, the motion of the aerial vehicle can be derived from the alignment of its images with the mosaic [3].

The main problem of this approach when large areas are covered is the small image motion inaccuracies, they lead to an erroneous estimation of the mosaic and, as a result, to a drift in the estimated position of the robot. In [2], an online mosaicking architecture that allows to improve the alignment of the central pixel of the images within the mosaic is presented; however, only translation errors are corrected. Pizarro [1] proposes an accurate method for image alignment and mosaic building, but its iterative nature leads to processing problems in online implementations.

This paper proposes a new technique in which a complete homography and its associated full covariance matrix are used to represent the alignment of the images in the mosaic. It draws an accurate model that considers the uncertainties in the local relations between the images that compose the mosaic. The proposed mosaic architecture allows the detection of loops and hence the reduction of the image drift which is present in this kind of methods. Furthermore, the uncertainty reduction in the local image alignment can be propagated to the rest of the mosaic thanks to the full covariance matrix computation.

This document describes the steps needed to introduce uncertainties in the mosaic building procedure by means of the Kalman filter. Firstly, an overview of the approach is presented. Then, the estimation of the local relations among images is outlined. Finally, the proposed Kalman filter for stochastic mosaicking is detailed and some experimental results are shown.

## II. AN OVERVIEW

The motivation of this work are the results presented in [3], where monocular imagery is used to compute the real motion that a camera attached to an aerial robot undergoes. The odometer presented there is based on the fact that it is possible to obtain the motion of a calibrated camera, up to a scale factor, from the homographic models that relate several images of a *planar* scene. The local homographies between consecutive images can be composed to obtain the global motion of the vehicle, but the local errors lead to a progressive drift in the estimated location of the vehicle. Furthermore, the mosaic can be used to reduce the cumulative errors associated to the odometry. However, the construction of the mosaic itself is affected by errors, which are not considered.

In the approach proposed in this paper, the mosaic building presented in [3] is improved at two levels. Firstly, a new approach is presented for the homography computation in pseudo-planar scenes where the planar assumption may not

This work is partially supported by AWARE project (IST-2006-33579) and AEROSENS project (DPI-2005-02293)

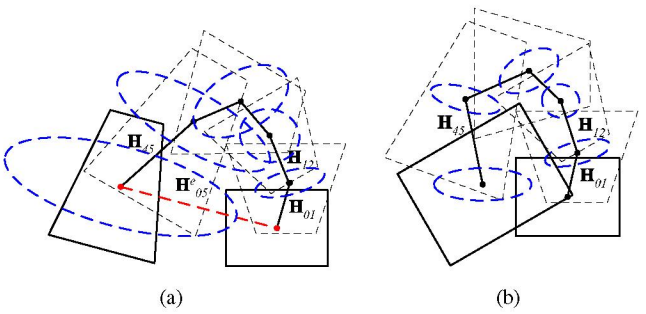


Fig. 1. Basic scheme: (a) when a part of the mosaic is revisited, (b) the drift is corrected and the correction propagated. Although only the covariances on the position are shown (which are used to determine potential loops), the full covariance matrix for all homographies is maintained.

hold. In addition, the covariance matrix associated to the homographies is computed in order to have an estimation of their accuracy.

Finally, the covariances of the estimated homographies are used in a new mosaic building architecture. When a zone of the mosaic is revisited, a procedure is used to reduce the accumulated drift. A full covariance matrix that considers the correlation among the estimated homographies in the loop is computed, and the stochastic information is employed to propagate the correction to all the images involved (see Fig. 1).

## III. HIERARCHICAL HOMOGRAPHY COMPUTATION

The technique proposed to build the mosaics is based on the homography computation between images. The algorithm used to compute the homography is described in detail in [4]. It basically consists of a point-feature tracker that obtains matches between images, and a combination of Least Median of Squares and a M-Estimator for outlier rejection and accurate homography estimation from these matches.

However, there are two factors that may reduce the applicability of the technique, mainly when the UAV flies at altitudes of the same order of other elements on the ground (buildings, trees, etc):

- In 3D scenes, the parallax effect will increase, and the planarity assumption may not hold. The result is a dramatic growth of the outliers and even the divergence of the M-Estimator.
- Depending on the frame-rate and the vehicle motion, the overlap between images in the sequence is sometimes reduced. This generates a non-uniform distribution of the features along the images.

The above problems are related, in the sense that both generate an ill-posed system of equations for the computation of the homography. If the matches are not uniformly distributed over the images, there may exist multiple solutions; and if the parallax effect is significant, there may exist multiple planes (whose transformation should be described by multiple homographies).

A classical solution to improve the results is to introduce additional constraints to reduce the number of degrees of freedom of the system of equations. In the proposed solution, this is accomplished through a *hierarchy* of homographic models (see Fig. 2), in which the complexity of the model to be fitted is decreased whenever the system of equations is ill-constrained.

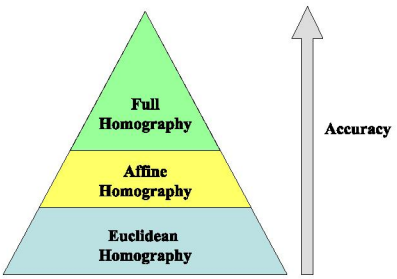


Fig. 2. Levels in the proposed hierarchical homography computation. Accuracy increases with the complexity of the model.

Therefore, depending on the quality of the available data, the constraints used to compute the homography are different; thus, the accuracy of the changes as well. An estimation of this accuracy will be given by the covariance matrix of the computed parameters. If this stochastic information is available, it can be used to define the accuracy of the image alignment inside the mosaic.

### A. Hierarchy implementation

The following paragraphs presents the way to progressively reduce the complexity of the homographic model and the laws that govern the transitions among the hierarchy.

Summarizing, any non-singular invertible $3x3$ matrix can be consider as homography:

$$\mathbf{H} = \begin{bmatrix} h_{00} & h_{01} & h_{02} \\ h_{10} & h_{11} & h_{12} \\ h_{20} & h_{21} & h_{22} \end{bmatrix} \quad (1)$$

A complete homography has 8 degrees of freedom (as it is defined up to a scale factor). The degrees of freedom can be reduced by fixing some of the parameters of the $3x3$ matrix. The models used are the defined by Hartley in [5]: Euclidean, Affine and Complete Homographic models, which have 4, 6 and 8 degrees of freedom respectively (see figure 2).

The percentage of successful matches obtained by the point tracker is used to have an estimation about the level of the hierarchy where the homography computation should start. These percentage thresholds were obtained empirically by processing hundreds of aerial images. Each level involves the following different steps:

- Complete homography. Least median of squares (LMedS) is used for outlier rejection and a M-Estimator to compute the final result. This model is used if more than the 65% of the matches are successfully tracked.
- Affine homography. If the percentage of success in the tracking step is between 40% and 65%, then the LMedS is not used, given the reduction in the number of matches. A relaxed M-Estimator (soft penalization) is carried out to compute the model.

- Euclidean homography. If the percentage is below 40%, the set of data is too noisy and small to apply non-linear minimizations. The model is computed using least-squares.

In addition, it is necessary a rule to know when the current level is not constrained enough and the algorithm has to decrease the model complexity. The M-Estimator used in the complete and affine computations is used for this. It is considered that the M-Estimator diverge if it reaches the maximum number of iterations and, hence, the level in the hierarchy has to be changed to the next one.

### B. Homography covariance estimation

Once the homography is computed, it is necessary to obtain a measure of the estimation accuracy. In this approach the covariance matrix is used as this accuracy measurement. By rearranging the homography matrix of (1) in the following vector:

$$\mathbf{h} = \begin{bmatrix} h_{00} & h_{01} & \dots & h_{22} \end{bmatrix}^T \tag{2}$$

the covariance is represented by a $9 \times 9$ matrix with the variance of these variables in the diagonal and the cross-variances in the upper and lower triangles.

Hartley [5] proposes a method to compute the covariance matrix of this estimation. Thus, given a set of $n$ matches:

$$\begin{aligned} S_m &= \{\{\mathbf{m}_1, \mathbf{m}'_1\}, \{\mathbf{m}_2, \mathbf{m}'_2\}, \dots \{\mathbf{m}_n, \mathbf{m}'_n\}\} \\ \mathbf{m}_i &= \begin{bmatrix} x_i \\ y_i \end{bmatrix}, \mathbf{m}'_i = \begin{bmatrix} x'_i \\ y'_i \end{bmatrix} \end{aligned} \tag{3}$$

The method can be summarized in the following 3 steps:

1) Obtain the Jacobian of the transformation from $\mathbf{m}$ to $\mathbf{m}'$ with respect to the 9 parameters of the homography $\mathbf{h}$, evaluated at the estimated homography by the previous procedure. This Jacobian is a $2n \times 9$ matrix derived from the homographic relation between the $n$ matches:

$$\begin{aligned} x'_i &= (h_{00}x_i + h_{01}y_i + h_{02})/(h_{20}x_i + h_{21}y_i + h_{22}) \\ y'_i &= (h_{10}x_i + h_{11}y_i + h_{12})/(h_{20}x_i + h_{21}y_i + h_{22}) \end{aligned} \tag{4}$$

   This Jacobian will be denoted as $\mathbf{J}$

2) Compute the covariance of the error for each match used to compute the homography (this is a $2x2$ matrix denoted by $\mathbf{C}_{\mathbf{m}_i}$). Assuming that the errors in the matches are uncorrelated, the diagonal matrix is created with the contribution of each match:

$$\mathbf{C}_{\mathbf{m}} = diag(\mathbf{C}_{\mathbf{m}_1}, \mathbf{C}_{\mathbf{m}_2}, \dots, \mathbf{C}_{\mathbf{m}_n}) \tag{5}$$

3) Known $\mathbf{J}$ and $\mathbf{C}_{\mathbf{m}}$, compute the homography covariance as:

$$\mathbf{C}_{\mathbf{h}} = (\mathbf{J}^T \mathbf{C}_{\mathbf{m}}{}^{-1} \mathbf{J})^{-1} \tag{6}$$

The unknown parameters in this algorithm are the error covariance for each match ($\mathbf{C}_{\mathbf{m}_i}$).

A rather straightforward approach could be to calculate the variance of the residue between $\mathbf{m}'_i$ and $\hat{\mathbf{m}}'_i = \mathbf{H}\mathbf{m}_i$ (the estimation of the match by means of the computed homography) and use it as error covariance for all the matches, so it is assumed that the error is approximately the same in the whole image. This method works well when the estimated homography is accurate, but usually fails when the estimation is poor, because the error may not be uniform.

Fig. 3. Image bucketing carried out to improve the covariance estimation.

To improve the calculation of the covariances a bucketing technique is proposed. The key idea is to divide the image in several sections (see figure 3) and compute the variance of the residue separately in each area. The bucketing selected takes into account the properties of the homographic transformations. In this way, each $\mathbf{C}_{\mathbf{m}_i}$ is the variance of the residue computed in the area to which the match $i$ belongs.

### C. Experimental results

This section shows some results of the proposed technique in real conditions. The algorithm was applied to a long image sequence captured at relatively low altitude, and the computed homographies were used to align each new image with a mosaic. This experiment tries to visually demonstrate the coherence of the homographies computed by using the proposed technique, not to build an accurate mosaic.

The images were taken by KARMA, an autonomous blimp developed by the RIA research team at LAAS, Toulouse [6]. The blimp is equipped with GPS, IMU and other navigation sensors like altimeter. In addition, it carries a digital camera and an on-board processor devoted to perception tasks.

Figure 4 shows the mosaic built with more than 400 images taken by KARMA. The scenario is the parking at LAAS, and the distance from the blimp to the ground is 27 meters. The low altitude makes evident the parallax effect; some trees are 15 meters high.

The technique described in [3] were applied to this image sequence but the parallax effect made impossible the homography computation during the processing and only a small portion of the mosaic were computed.

The image sequence used to build the mosaic (figure 4) was processed at a rate of 2 images per second with resolution of 1024x768. The final mosaic has a 4300x2400 effective resolution.

## IV. STOCHASTIC MOSAICKING

The goal of this section is to improve the environment model by attaching stochastic information to the mosaic. This information will be used to improve the positions of

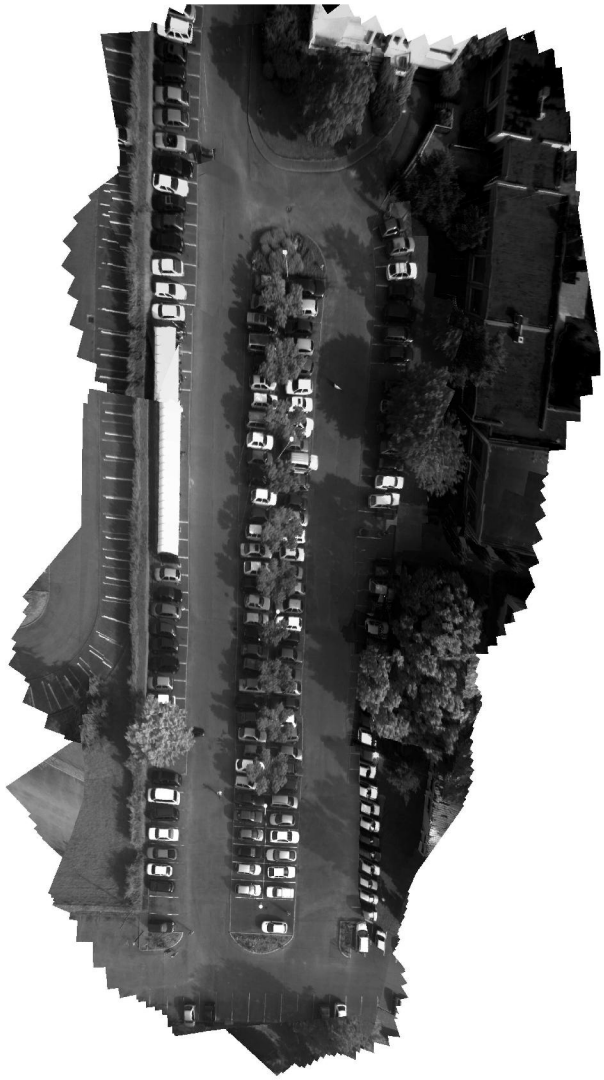

Fig. 4. Mosaic of "le petit parking" at LAAS (Toulouse). Images taken at 27 meter over ground. Images taken by KARMA

the images within the mosaic when a close-loop is detected. In this approach, the mosaic is represented by a database of images and associated data; the mosaic is no longer a static image, but a set of images linked by stochastic relations.

When a close-loop is detected the relations among the images involved in the loop are updated according to the accuracy of the measurement and the estimation of the position of each image. A minimization process based on a Kalman Filter is used to optimize the estimation of the inter-image relations.

### A. Closing the loop

Each time a new image is gathered by the cameras of the UAV, the homography that relates it with the previous one is computed, as well as its covariance. The position of the image inside the mosaic is obtained multiplying the current homography by all the previous homographies until the reference frame is reached.

The loop-closing is computed by detecting the crossover of the current image with one or more images of the database. For this purpose, the estimated position (and its covariance) of the central pixel of each image in the mosaic is stored in the database. The crossover detection consists of finding one image whose Mahalanobis distance to its central pixel is within a certain empirical range.

Once the crossover is detected, a feature matching procedure is launched to compute the alignment between both images. In the general case, the task of matching images taken from very different points of view is difficult and computationally costly. Even if a good estimation of the translation between images is available, the tracking algorithm [4] cannot deal with a rotation higher that 45 degrees, as the features used are not affine invariant.

However, due to the pseudo-planar nature of the scene, the matching complexity can be drastically reduced. An initial estimation of the location (a complete homography) of the image inside the mosaic is available, so this estimation can be used as a searching seed for the feature matching procedure. At time $i$ the homography that aligns the current image $I_i$ to the mosaic is $\hat{\mathbf{H}}_{0i}$. If a crossover with image $I_j$, whose estimated position inside the mosaic is given by $\hat{\mathbf{H}}_{0j}$, is detected, the initial estimation of the transformation from $I_i$ to $I_j$ is defined by:

$$\hat{\mathbf{H}}_{ji} = \prod_{k=j+1}^{i} \hat{\mathbf{H}}_{(k-1)k} \tag{7}$$

If the estimation error for each homography were zero, the alignment of $I_i^w$ (the result of warping $I_i$ according to $\hat{\mathbf{H}}_{ji}$) with respect $I_j$ would be perfect, but inaccuracies in the estimated homographies produce unavoidable alignment errors.

However, computing matches between the warped image $I_i^w$ and $I_j$ is an easier task. From this data is possible to obtain the homography $\mathbf{H}_{ji}^e$ that describes the drift between both images and then obtain the correct alignment:

$$\mathbf{H}_{ji} = \hat{\mathbf{H}}_{ji}\mathbf{H}_{ji}^e \tag{8}$$

It is clear that by using this method the matching algorithm does not have to deal with the complete motion between both images, only with the accumulated errors. It will reduce the complexity of the matching algorithm needed and will improve the alignment results.

In addition, using normalized correlation allows to deal with the illumination problems that appears in close loops when images are taken at different time of day. The erroneous matches associated to different causes like shadows and parallax are later removed by the outlier detection and robust estimation performed in the homography computation.

### B. Updating the mosaic after a loop-closing

As stated before, when a loop-closing is detected a measurement about the alignment between the current image $I_i$ and the crossover $I_j$ is given by the system. This measure is a homography $\mathbf{H}_{ji}$ and its covariance matrix $\mathbf{C}_{\mathbf{h}_{ji}}$. The problem

now is how to update the relations among the images from $I_j$ to $I_i$ under the following constraint:

$$\mathbf{H}_{ji} = \prod_{k=j+1}^{i} \hat{\mathbf{H}}_{(k-1)k} \tag{9}$$

For this purpose, a minimization process based on the Extended Kalman Filter is launched each time a loop-closing is detected. This filter re-estimates the relations among the images involved in the loop, taking into account the uncertainty of the loop-closing and the stochastic information stored in the database. The next sections will detail the structure and dynamics of the filter.

*1) The state vector:* It is assumed that a close-loop is detected between $I_j$ and $I_i$, with $n+1$ images involved, and the transformation that aligns both crossover images is $\mathbf{H}_{ji}$ with covariance $\mathbf{C}_{\mathbf{h}_{ji}}$. For simplicity, it can be supposed that $j=1$ and $i=n+1$.

The a priori state vector will be composed by the $n$ transformations that align the $n$ images $(I_2,...,I_{n+1})$ with $I_1$, thus:

$$\mathbf{x}^- = [\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n]^T = [\mathbf{h}_{12}, \mathbf{x}_1 \cdot \mathbf{h}_{23}, ..., \mathbf{x}_{n-1} \cdot \mathbf{h}_{n(n+1)}]^T \tag{10}$$

where $\mathbf{x}_i$ and $\mathbf{h}_{ij}$ are the vector representation of the homography matrices $\mathbf{X}_i$ and $\mathbf{H}_{ij}$ , as shown in equation (2). The operator "$\cdot$" is the product of two homographic vectors, and the result is expressed as a vector.

It is easy to see that the state is obtained by recursively applying equation (7): $\mathbf{x}_1$ is the product at $k=j+1=2$, $\mathbf{x}_2$ is the product at $k=j+2=3$ that can be written in function of $\mathbf{x}_1$ and so on.

The calculation of the a priori state vector and its a priori covariance matrix $\mathbf{P}^-$ is done in an incremental fashion for the $n$ states using the following prediction equation:

$$\mathbf{x}_i = \mathbf{x}_{i-1} \cdot \mathbf{h}_{i(i+1)} \tag{11}$$

It is easy to compute the Jacobian of this expression with respect to the state vector (matrix $\mathbf{A}$) and with respect to the variables of $\mathbf{h}_{i(i+1)}$ (matrix $\mathbf{W}$). $\mathbf{A}$ will be used to generate the corresponding rows and columns of the matrix $\mathbf{P}^-$ (see figure 5) and, with $\mathbf{W}$, to compute the covariance matrix of $\mathbf{x}_i$ by means of:

$$\mathbf{C}_{\mathbf{x}_i} = \mathbf{A}\mathbf{C}_{\mathbf{x}_{i-1}}\mathbf{A}^T + \mathbf{W}\mathbf{C}_{\mathbf{h}_{i(i+1)}}\mathbf{W}^T \tag{12}$$

The structure of the full covariance matrix shows that every homography is correlated with the previous ones in the loop.

*2) Updating the mosaic:* The previously described state vector is arranged in order to minimize the number of operations needed to update it with the measurement $(\mathbf{H}_{1(n+1)}, \mathbf{C}_{\mathbf{h}_{1(n+1)}})$. Thus, the state $\mathbf{x}_n$ of $\mathbf{x}^-$ represents the transformation from the current image $n+1$ to the loop reference frame 1, so $\mathbf{h}_{1(n+1)} = \mathbf{x}_n$. This fact is described in the following measurement equation:

$$\mathbf{h}_{1(n+1)} = \mathbf{F}\mathbf{x} = \left[\mathbf{0}_{9x(9n)}, \mathbf{I}_{9x9}\right]\mathbf{x} \tag{13}$$

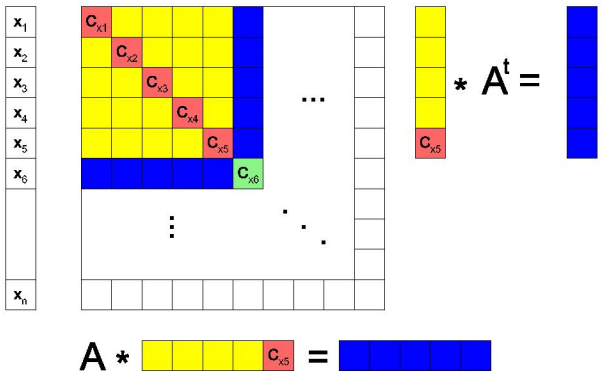


Fig. 5. Generation of the a priori covariance matrix of the state vector

It is easy to update the state vector and its covariance matrix following the classic Extended Kalman Filter equations. This simple measurement function $\mathbf{F}$ will significantly reduce the computation needed to obtain the gain $\mathbf{K}$ and, later, the a posteriori covariance matrix $\mathbf{P}$.

Special attention is required for the state vector updating:

$$\mathbf{x} = \mathbf{x}^- + \mathbf{K}(\mathbf{h}_{1(n+1)} - \mathbf{F}\mathbf{x}^-) \tag{14}$$

The homographies are defined up to scale factor; thus, a homography multiplied by a constant $k$ represents exactly the same transformation, although its components are different to those of the original matrix. To implement the substraction of homographies in (14) it is necessary to perform a normalization of the scale factor. The proposed solution is to set the determinant of both homographies to 1 before computing the state vector.

Finally, the relations among the images involved in the loop are updated with the measurement but they are expressed in the local coordinate frame of the loop (the first image of the loop, $I_1$, is the reference frame) and have to be transformed into the mosaic system reference. From equation (11) can be easily obtained:

$$\mathbf{H}_{i(i+1)} = (\mathbf{X}_{i-1})^{-1}\mathbf{X}_i \tag{15}$$

and the covariances can be derived from the Jacobian of this expression.

## V. Results of Vision-Based Localization Experiments

This section shows some experiments to demonstrate the correct operation of the proposed algorithm. The framework of the experiments is the position computation by means of monocular imagery. The position of the vehicle is estimated by means of the algorithm described in [3] and the inputs are the homographies computed with the proposed stochastic mosaicking.

The images were taken by KARMA during one experiment in the parking at LAAS. The sequence is composed by two hundred images. Figure 6 shows the mosaic built with the images and the computed homographies.

More that 15 short loops (less than 20 images involved) and one large loop (more than 100 images) were closed

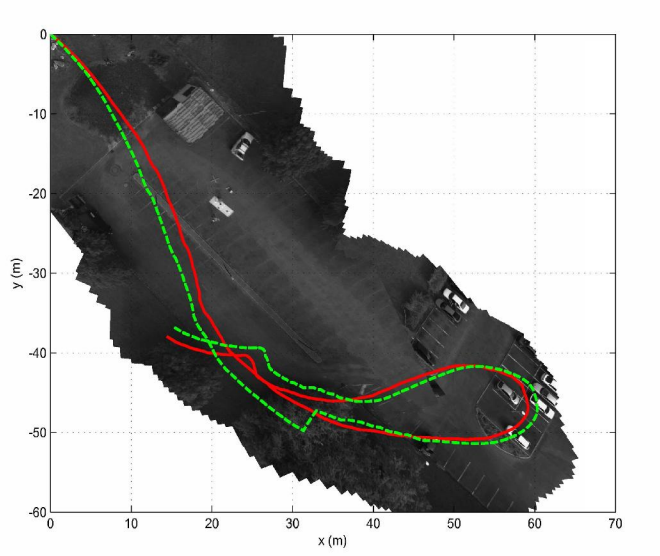


Fig. 6. Position estimation experiment with mosaic computation. Mosaic of "le grand parking" at LAAS (Toulouse). KARMA was flying at 22 meter over the ground. Image based XY position estimation (dashed line) and GPS position measurement (solid line) are displayed.

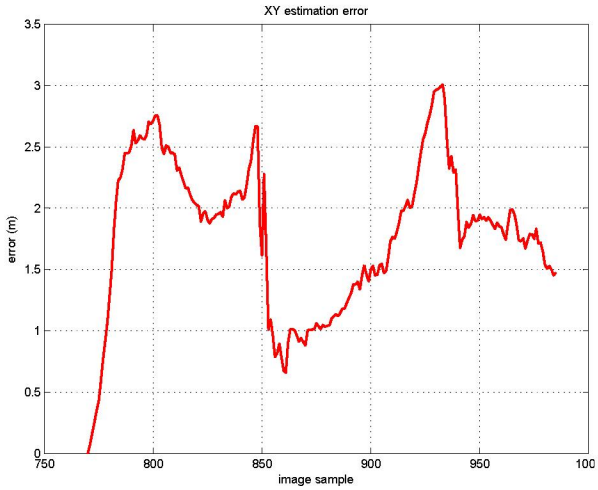


Fig. 7. XY estimation error per image sample.

during the experiment. In figure 6 the image based XY position estimation and GPS position measurement are displayed. Notice the sudden left turn in the vision based estimation at coordinates (31, -50), it is automatically carried out when the plane inducted by the floor (our ground truth) becomes the scene dominant plane again.

Figure 7 shows the evolution of the errors associated to the position estimation. The error is calculated as the Euclidean distance between the GPS and the estimated position at each image sample. It can be seen that the error is moderate: its maximum value is 3 meters and the mean is 1.76 meters.

## VI. CONCLUSIONS AND FUTURE WORKS

Mosaics can be a convenient environment representation for UAV tasks, such as monitoring events, identification of changes and others. The paper presents an approach for building mosaics based on planar homographies. A hierarchical homography computation increases the robustness of the approach, allowing to deal with quasi-planar scenes.

An important aspect of the paper is the inclusion of uncertainty measures in the mosaic. The mosaic is no longer a static picture, but a collection of relations among images. Although the alignment errors will grow over time, they can be partially reset when part of the mosaic is revisited. Moreover, the correction can be propagated to the rest of the mosaic by considering the relations among images.

UAV position estimation is an interesting potential application of the proposed approach. It is possible to determine the motion of a calibrated camera [3] from plane-induced homographies. The estimated motion is affected by the drift of the homographies computed. The current procedure allows to reduce it.

The main drawback of the approach is the storage requirements. Not only the images of the mosaic should be stored (although this number can be controlled, depending on the degree of overlap), but also the state augments along the sequence of images. The current state of the technology allows to solve the image storage by means of solid state hard disks which support high capacity at low weight. However, this approach can't deal in real-time with the computation needed to update the mosaic in very large loops. In this sense, new techniques based on the downsampling of the homographic relations between images and the use of Sequential Map Joining [7] will be researched.

## VII. ACKNOWLEDGMENTS

The authors thank to the Laboratoire d'Architecture et d'Analyse des Systèmes at Toulouse, France, for supporting the images and telemetry of KARMA used to validate the algorithms. Special thanks to Simon Lacroix, Sebastien Bosch and Tomas Lemaire for their help and good advices during the development of this work.

## REFERENCES

[1] O. Pizarro and H. Singh, "Toward large-area mosaicing for underwater scientific applications," *IEEE Journal of Oceanic Engineering*, vol. 28, no. 4, pp. 651–672, October 2003.

[2] R. Garcia, J. Puig, P. Ridao, and X. Cufi, "Augmented state kalman filtering for auv navigation," in *Proceedings of the International Conference on Robotics and Automation*, Washington, 2002., pp. 4010–4015.

[3] F. Caballero, L. Merino, J. Ferruz, and A. Ollero, "Improving vision-based planar motion estimation for unmanned aerial vehicles through online mosaicing," in *Proceedings of the International Conference on Robotics and Automation*. IEEE, May 2006.

[4] A. Ollero, J. Ferruz, F. Caballero, S. Hurtado, and L. Merino, "Motion compensation and object detection for autonomous helicopter visual navigation in the comets system," in *Proceedings of the International Conference on Robotics and Automation, ICRA*. IEEE, 2004, pp. 19–24.

[5] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, 2004.

[6] E. Hygounenc, I.-K. Jung, P. Soueres, and S. Lacroix, "The Autonomous Blimp Project of LAAS-CNRS: Achievements in Flight Control and Terrain Mapping," *The International Journal of Robotics Research*, vol. 23, no. 4-5, pp. 473–511, 2004.

[7] J. Tardos, J. Neira, P. Newman, and J. Leonard, "Robust mapping and localization in indoor environments using sonar data," *Int. J. of Robotics Research*, vol. 24, no. 4, pp. 311–330, 2002.